

04-27-00

A

04/26/00



jc808 U.S. PTO

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

jc525 U.S. PTO

09/560269



04/26/00

In re Patent Application of: Barry M. Nolte

Title: MINIMAL INSTRUMENTATION FOR LOSSLESS CALL PROFILING

Attorney Docket No.: 777.344US1

PATENT APPLICATION TRANSMITTAL


BOX PATENT APPLICATIONAssistant Commissioner for Patents
Washington, D.C. 20231

We are transmitting herewith the following attached items and information (as indicated with an "X"):

- X Utility Patent Application under 37 CFR § 1.53(b) comprising:
- X Specification (25 pgs, including claims numbered 1 through 45 and a 1 page Abstract).
 - X Formal Drawing(s) (5 sheets).
 - X Unsigned Combined Declaration and Power of Attorney (3 pgs).
 - X Return postcard.

The filing fee (NOT ENCLOSED) will be calculated as follows:

	No. Filed	No. Extra	Rate	Fee
TOTAL CLAIMS	45 - 20 =	25	x 18 =	\$450.00
INDEPENDENT CLAIMS	6 - 3 =	3	x 78 =	\$234.00
1 MULTIPLE DEPENDENT CLAIMS PRESENTED				\$0.00
BASIC FEE				\$690.00
TOTAL				\$1,374.00


THE FILING FEE WILL BE PAID UPON RECEIPT OF THE NOTICE TO FILE MISSING PARTS.SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.
P.O. Box 2938, Minneapolis, MN 55402 (612-373-6900)By: 
Atty: Allen J. Oh
Reg. No. 42,047**Customer Number 21186**

"Express Mail" mailing label number: EL510128190US.

Date of Deposit: April 26, 2000

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Box Patent Application, Washington, D.C. 20231.

By: Shawn L. Hise

Signature: 

MINIMAL INSTRUMENTATION FOR LOSSLESS CALL PROFILING

Copyright Notice and Permission

A portion of the disclosure of this patent document may contain material that is
5 subject to copyright protection. The copyright owner has no objection to the facsimile
reproduction by anyone of the patent document or the patent disclosure, as it appears in
the Patent and Trademark Office patent files or records, but otherwise reserves all
copyright rights whatsoever. The following notice shall apply to this document:
Copyright © 2000, Microsoft, Inc.

Technical Field of the Invention

The present invention pertains generally to development of computer software
applications, and more particularly to performance analysis of such applications.

Background of the Invention

Computer software applications typically consist of many modules, known as
functions, which work together to perform desired tasks. A significant concern in
software development is operational speed, which is largely dependent on the
performance of the individual functions making up the application. Thus, an important
20 component of application development is monitoring application performance.

To provide useful information, performance measurement should, to the extent
possible, avoid interfering with the normal operation of the application. Because
performance measurement involves collecting data in circumstances under which data
would not ordinarily be collected, performance measurement necessarily affects the
25 normal operation of the application to a certain extent. Certain factors contribute to this
effect. For example, inserting data collection code in an application makes the
application larger, which slows operation down somewhat. This decrease in performance
is attributable to the increase in the amount of code that must be both loaded into the
system's memory and executed. In addition, adding data collection code increases the
30 amount of time other operations take due to the execution of the data collection code. As
a result, the time relationships between the other activities within the application and

within the system are affected. Data collection also affects application performance as a result of input/output (I/O) bandwidth considerations. Specifically, when data is collected, it must be stored, for example, in physical memory, on a disk, or on a networked device. Storing data consumes significant time due to the relatively slow process of outputting data to the storage device. This burden is especially considerable when large amounts of data are collected. While the burden can be reduced by adding a dedicated I/O device, such as a disk drive, device contention may still arise at the system level.

Summary of the Invention

According to various example implementations of the invention, there is provided a process for collecting information relating to the execution of an application for use in performance analysis. Probes are inserted at only those points in the application that would yield non-redundant information. By avoiding the collection of redundant information, the effect of the information collection process on the application is reduced. Additionally, performance of the data analysis process is improved due to the reduction of cases in which it must be decided whether data is redundant, and also due to the reduction of work involved in discarding redundant data. The accuracy and usefulness of the collected information are thus improved.

In a particular implementation, probe locations are determined at which collecting data relating to the execution of the application would produce non-redundant information. Probes are then inserted only at these locations. In general, these locations include the entry and exit points of the functions making up the application. Other particular implementations of the present invention relate to, for example, exceptional cases in which probes are inserted at other locations in addition to entry and exit points.

In another particular implementation, the entry and exit points of functions constituting at least a portion of the application are determined. Also determined are probe locations in the application at which collecting data relating to the execution of the application would produce non-redundant information. These probe locations include at least the determined entry and exit points. Probes are inserted only at the determined

probe locations in the application. These probes are then used to collect the information relating to the execution of the application. The collected information is then analyzed.

Still other implementations include computer-readable media and apparatuses for performing these methods. The above summary of the present invention is not intended to describe every implementation of the present invention. The figures and the detailed description that follow more particularly exemplify these implementations.

Brief Description of the Drawings

- Figure 1 illustrates a simplified overview of an example embodiment of a computing environment for the present invention.
- Figure 2 is a flow chart illustrating an example method for analyzing the performance of an application, according to an embodiment of the present invention.
- Figure 3 depicts a portion of an application with potential probe locations.
- Figure 4 depicts an example instrumentation scheme according to one operational example of an embodiment of the present invention.
- Figure 5 depicts another example instrumentation scheme according to another operational example of an embodiment of the present invention.

Detailed Description of the Invention

In the following detailed description of the preferred embodiments, reference is made to the accompanying drawings that form a part hereof, and in which are shown by way of illustration specific embodiments in which the invention may be practiced. It is understood that other embodiments may be utilized and structural changes may be made without departing from the scope of the present invention.

Hardware and Operating Environment

Figure 1 illustrates a hardware and operating environment in conjunction with which embodiments of the invention may be practiced. The description of Figure 1 is intended to provide a brief, general description of suitable computer hardware and a

suitable computing environment with which the invention may be implemented.

Although not required, the invention is described in the general context of computer-executable instructions, such as program modules, being executed by a computer, such as a personal computer (PC). This is one embodiment of many different computer configurations, some including specialized hardware circuits to analyze performance, that may be used to implement the present invention. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types.

Moreover, those skilled in the art will appreciate that the invention may be practiced with other computer-system configurations, including hand-held devices, multiprocessor systems, microprocessor-based or programmable consumer electronics, network personal computers ("PCs"), minicomputers, mainframe computers, and the like. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

Figure 1 shows a computer arrangement implemented as a general-purpose computing or information-handling system 80. This embodiment includes a general purpose computing device such as personal computer (PC) 120, that includes processing unit 121, a system memory 122, and a system bus 123 that operatively couples the system memory 122 and other system components to processing unit 121. There may be only one or there may be more than one processing unit 121, such that the processor computer 120 comprises a single central-processing unit (CPU), or a plurality of processing units, commonly referred to as a parallel processing environment. The computer 120 may be a conventional computer, a distributed computer, or any other type of computer; the invention is not so limited.

In other embodiments other configurations are used in PC 120. System bus 123 may be any of several types, including a memory bus or memory controller, a peripheral bus, and a local bus, and may use any of a variety of bus architectures. The system memory 122 may also be referred to as simply the memory, and it includes read-only memory (ROM) 124 and random-access memory (RAM) 125. A basic input/output

system (BIOS) 126, stored in ROM 124, contains the basic routines that transfer information between components of personal computer 120. BIOS 126 also contains start-up routines for the system.

Personal computer 120 further includes hard disk drive 127 having one or more magnetic hard disks (not shown) onto which data is stored and retrieved for reading from and writing to hard-disk-drive interface 132, magnetic disk drive 128 for reading from and writing to a removable magnetic disk 129, and optical disk drive 130 for reading from and/or writing to a removable optical disk 131 such as a CD-ROM, DVD or other optical medium. Hard disk drive 127, magnetic disk drive 128, and optical disk drive 130 are connected to system bus 123 by a hard-disk drive interface 132, a magnetic-disk drive interface 133, and an optical-drive interface 134, respectively. The drives 127, 128, and 130 and their associated computer-readable media 129, 131 provide nonvolatile storage of computer-readable instructions, data structures, program modules and other data for personal computer 120. Although the exemplary environment described herein employs a hard disk, a removable magnetic disk 129 and a removable optical disk 131, those skilled in the art will appreciate that other types of computer-readable media which can store data accessible by a computer may also be used in the exemplary operating environment. Such media may include magnetic tape cassettes, flash-memory cards, digital video disks (DVD), Bernoulli cartridges, RAMs, ROMs, and the like.

In various embodiments, program modules are stored on the hard disk drive 127, magnetic disk 129, optical disk 131, ROM 124 and/or RAM 125 and may be moved among these devices, *e.g.*, from hard disk drive 127 to RAM 125. Program modules include operating system 135, one or more application programs 136, other program modules 137, and/or program data 138. A user may enter commands and information into personal computer 120 through input devices such as a keyboard 140 and a pointing device 142. Other input devices (not shown) for various embodiments include one or more devices selected from a microphone, joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 121 through a serial-port interface 146 coupled to system bus 123, but in other embodiments they are connected through other interfaces not shown in Figure 1, such as a parallel port, a game port, or a universal serial bus (USB) interface. A monitor 147 or other display

device also connects to system bus 123 via an interface such as a video adapter 148. In some embodiments, one or more speakers 157 or other audio output transducers are driven by sound adapter 156 connected to system bus 123. In some embodiments, in addition to the monitor 147, system 180 includes other peripheral output devices (not shown) such as a printer or the like.

In some embodiments, personal computer 120 operates in a networked environment using logical connections to one or more remote computers such as remote computer 149. Remote computer 149 may be another personal computer, a server, a router, a network PC, a peer device, or other common network node. Remote computer 149 typically includes many or all of the components described above in connection with personal computer 120; however, only a storage device 150 is illustrated in Figure 1. The logical connections depicted in Figure 1 include local-area network (LAN) 151 and a wide-area network (WAN) 152, both of which are shown connecting PC 120 to remote computer 149; typical embodiments would only include one or the other. Such networking environments are commonplace in offices, enterprise-wide computer networks, Intranets and the Internet.

When placed in a LAN networking environment, PC 120 connects to local network 151 through a network interface or adapter 153. When used in a WAN networking environment such as the Internet, PC 120 typically includes modem 154 or other means for establishing communications over network 152. Modem 154 may be internal or external to PC 120 and connects to system bus 123 via serial-port interface 146 in the embodiment shown. In a networked environment, program modules depicted as residing within PC 120 or portions thereof may be stored in remote-storage device 150. Of course, the network connections shown are illustrative, and other means of establishing a communications link between the computers may be substituted.

Software may be designed using many different methods, including object-oriented programming methods. C++ and Java are two examples of common object-oriented computer programming languages that provide functionality associated with object-oriented programming. Object-oriented programming methods provide a means to encapsulate data members (variables) and member functions (methods) that operate on

that data into a single entity called a class. Object-oriented programming methods also provide a means to create new classes based on existing classes.

An object is an instance of a class. The data members of an object are attributes that are stored inside the computer memory, and the methods are executable computer code that act upon this data, along with potentially providing other services. The notion of an object is exploited in the present invention in that certain aspects of the invention are implemented as objects in some embodiments.

An interface is a group of related functions that are organized into a named unit. Some identifier may uniquely identify each interface. Interfaces have no instantiation; that is, an interface is a definition only without the executable code needed to implement the methods that are specified by the interface. An object may support an interface by providing executable code for the methods specified by the interface. The executable code supplied by the object must comply with the definitions specified by the interface. The object may also provide additional methods. Those skilled in the art will recognize that interfaces are not limited to use in or by an object-oriented programming environment.

Example Embodiments of the Invention

According to one example embodiment of the present invention, probes are intelligently inserted in applications to collect enough information to perform full call attribution, *i.e.*, to measure the amount of time any function takes from start to finish, while avoiding collecting redundant information that is not needed to perform full call attribution. By avoiding the collection of unnecessary information, the various implementations of the present invention reduce the effect of the instrumentation process on the application being instrumented. As a result, normal operations are disrupted less significantly, allowing for greater accuracy in measuring application performance. Instrumentation can be performed at compilation time, as a post link step, or even while the application is running. Furthermore, the techniques described herein are language-independent and can be used with all procedural languages, *e.g.*, C, C++, Java, Pascal, Fortran, BASIC, Cobol, *etc.*

In particular, the application whose performance is to be measured is instrumented at logical points that can be associated with the original source code that

was used to generate the application, while maintaining a granularity that does not generate data that is not needed to perform full call attribution. In most cases, probes are inserted at the entry and exit points of all discrete functions in the application. Under certain circumstances, however, probes are also placed at other strategic points in the application. Intelligent placement of these probes ensures that enough information is collected to perform full call attribution, while avoiding unduly interfering with the normal operation of the application by avoiding the collection of redundant information.

According to a particular implementation of the present invention, various types of probes are used to monitor the performance of an application. One type labeled, for example, `_EnterFunction`, is placed at the entry point of a function and collects the address of the function being entered and the current stack pointer, as well as a time stamp that indicates when the function was entered. Another type of probe, an `_ExitFunction` probe, is placed at an exit point of a function and collects data similar to that collected by the `_EnterFunction` probe. In particular, the `_ExitFunction` probe collects the address of the function being exited, the current stack pointer, and a time stamp that indicates when the function was exited. Yet another type of probe, a `_CallToFunction` probe, is placed at certain points in the application immediately before program flow changes from one function to another. This type of probe collects the address of the function from which the call originated, the address of the function being called, the current stack pointer, and a time stamp indicating when the change in program flow occurred. A fourth type of probe, a `_ReturnFromFunction` probe, is placed at the return point of a program flow change. The `_ReturnFromFunction` probe collects the address of the function to which program flow returns, the current stack pointer, and a time stamp indicating when program flow returned to the previous function.

For each of these types of probes, the address of functions is collected so that, when performance is analyzed, time can be accurately attributed to the correct functions. The probes also collect the current stack pointer to facilitate correct attribution of recursive function calls. That is, collecting the current stack pointer allows the correct attribution of time to the correct instance of a function in recursive cases. In addition, attribution is facilitated in certain cases in which program flow is changed in non-traditional ways.

Referring again to the drawings, Figure 2 depicts an example method 200 for analyzing application performance, according to a particular embodiment of the present invention. At a block 202, the functions are analyzed to determine their entry and exit points. In particular, a function's entry point is the first statement in the function. Its exit points typically include the last statement in the function, as well as any points within the function at which other functions are called.

Each of these entry and exit points is a potential location at which a probe might be inserted to collect data relevant to full call attribution. While full call attribution can be performed by inserting probes at all of these locations, it is usually not necessary to collect data at all of these points. Indeed, omitting instrumentation at certain non-essential points reduces the effect of the performance analysis process on the application whose performance is being analyzed, thereby improving accuracy. Accordingly, at a block 204, the entry and exit points of the function are analyzed to determine which points, if instrumented with a probe, would yield non-redundant information, *i.e.*, which points should be instrumented.

Figure 3 depicts a portion 300 of an application and its potential probe locations 302, 304, 306, 308, 310, and 312. For purposes of clarity, the potential probe locations are depicted in italic type in Figure 3. As can be seen in Figure 3, potential locations for the `_EnterFunction` probes 302 and 310 occur at the beginnings of a main function 314 and of a `CalledToFunction` function 316. Similarly, potential locations for the `_ExitFunction` probes 308 and 312 occur at the ends of the main function 314 and of the `CalledToFunction` function 316.

Some of the probes can be omitted so as to collect sufficient information to perform full call attribution without collecting redundant information. At the block 204 of Figure 2, the system determines which probes can be omitted. In the specific example shown in Figure 3, the `_CallToFunction` and `_ReturnFromFunction` probes provide redundant information. These probes are redundant because they would supply almost exactly the same information at almost the same time as other probes. For example, the `_CallToFunction` probe 304 in the main function 314 would supply almost the same information as the `_EnterFunction` probe 310 in the function 316 at nearly the same time, with only a slight time difference attributable to overhead for the probe itself. Similarly,

the `_ReturnFromFunction` probe 306 in the main function 314 would supply almost the same information as the `_ExitFunction` probe 312 in the function 316 at nearly the same time. Thus, because the `_CallToFunction` and `_ReturnFromFunction` probes are redundant in this example, the potential probe locations 304 and 306 need not be instrumented, and are eliminated from consideration by the system at the block 204 of Figure 2. It is desirable to eliminate the `_CallToFunction` and `_ReturnFromFunction` probes because there are typically more locations that call to functions than there are functions themselves. Thus, eliminating these probe pairs removes a larger portion of the total instrumentation, while ensuring that enough information to perform full call attribution is collected.

While this strategy increases efficiency in most cases, it does not work for all cases. In particular, eliminating these probe pairs only works in the case of direct calls to functions within the current module. Thus, while instrumentation of calls within the current module can be eliminated from the list of instrumentation points, instrumentation of calls outside of the current module cannot. Eliminating these probes would cause insufficient information for full call attribution to be gathered, since there is no guarantee that other modules will be instrumented. Time can be incorrectly attributed as a result. For example, referring again to Figure 3, if the function 316 were located in a different module from the main function 314, eliminating the probes at locations 304 and 306 would cause time spent in the function 316 to be incorrectly attributed to the main function 314. Thus, as a general rule, calls outside the current module are instrumented both before and after the call is performed, *e.g.*, if the main function 314 and the function 316 were located in different modules, `_CallToFunction` and `_ReturnFromFunction` probes would be inserted at locations 304 and 306, respectively.

Another situation in which the normal strategy of instrumenting only the entry and exit points of functions produces potentially inaccurate results occurs in what is known as a tail merge. A tail merge is an optimization made either by a compiler or by the programmer at the machine instruction level that eliminates calls to functions by instead jumping to the next function, then having a single return statement return out of multiple functions. Figure 4 depicts an example of a tail merge situation. In Figure 4, a main function 402 calls a Function1 function 404. At a point 406, the function 404 in

turn calls a Function2 function 408. At the end of this function, a return statement 410 is encountered. This return statement 410, however, returns to the main function 402. This change in flow is contrary to the expected sequence of events, in which the return statement 410 would return to the function 404, which called the function 408.

- 5 Normally, another return statement would redirect execution from the function 404 back to the main function 402, which called the function 404. Because a tail merge functions similarly to a call/return pair, it should be instrumented. A `_CallToFunction` probe 412 is inserted before the jump at the point 406, treating the jump as a call to the function 408. Since execution does not return to the point 406 after the jump, however, a
- 10 `_ReturnFromFunction` probe placed after the jump would be superfluous.

- 15 This departure from the normal sequence of events is a potential source of ambiguity in attributing time to the correct functions. Specifically, an issue arises as to the proper attribution of a time period 414 occurring after the point 406 and before the end of the main function 402. Typically, an `_ExitFunction` probe would precede the return statement 410 and would indicate that program flow is exiting the function 408. This probe, however, does not convey any information as to where program flow continues after exiting the function 408. Without additional information, this time is likely to be erroneously attributed to the function 404 because the function 408 was called by the function 404, but no end point was available within the function 404 with
- 20 which to properly attribute time. The time period 414, however, should be attributed to the main function 402, which is where execution proceeds after the jump. This ambiguity can be resolved in a number of ways. One way is to flag the `_CallToFunction` probe 412 as calling a function that has a tail merge. Another way is to insert a `_ReturnFromFunction` probe 416 at the point to which execution proceeds after the tail
- 25 merge. Inserting the `_ReturnFromFunction` probe 416 at this location clarifies the destination of the return from the function 408 when the function trace is processed.

- 30 Another scenario in which inserting probes at only the entry and exit points of functions does not provide sufficient information to perform full call attribution arises in connection with blocks of code known as catch blocks and in connection with non-local go-to's. Both of these situations arise most commonly in connection with error handling mechanisms and can be instrumented in the same manner. Figure 5 illustrates a catch

block and its instrumentation. In the example depicted in Figure 5, a block of code attached to a catch statement 502, known as a catch block 502, is not called. Rather, execution jumps to the catch block 502 in the case of an exception throw 506 in a try block 508. Execution of the catch block 502 is neither a function call nor a tail merge.

- 5 The catch block 502, however, can be treated as a function for purposes of performance analysis. In one implementation, the catch block 502 is treated as a reentrant call to the function that contains the catch block 502, *i.e.*, a main function 510 in this case. Thus, in this example, an `_EnterFunction` probe 512 and an `_ExitFunction` probe 514 are inserted in the catch block 502 at the locations illustrated in Figure 5, with the function being
- 10 entered and exited denoted as the main function 510.

- Some other error handling mechanisms produce non-local go-to's by causing execution to jump to other blocks of code. For example, in the C programming language, the `setjmp/longjmp` facilities produce non-local go-to's. These blocks of code can be instrumented in the same manner as catch blocks, as discussed above in connection with
- 15 Figure 5.

- Another situation that is noteworthy is that raised by inlined functions, which are functions marked as "inlined" by the developer as a directive to the compiler. A function is requested to be inlined when the developer surmises that the cost (in terms of the code generated) associated with calling, executing, and returning from the function is greater
- 20 than the cost associated with simply executing the code "inline" within the calling function. The compiler may or may not fulfill this request. If the request is fulfilled, the inlined function will become part of another function and will not have entry and exit points associated with it. Thus, the inlined function is not instrumented. If, on the other hand, the request is not fulfilled, then the function that was intended to be inlined will
- 25 still have entry and exit points, which are instrumented with `_EnterFunction` and `_ExitFunction` probes, as appropriate. These probes are useful in that the developer does not expect the function to exist as a distinct entity. If time is attributed to a function that the developer expects to be inlined, then the developer can investigate the circumstances precluding the inlining of the function and can make changes to enable the function to be
- 30 inlined.

Referring again to Figure 2, after the non-redundant probe locations are determined in connection with the block 204, execution proceeds to a block 206, at which the appropriate probes are inserted in the application at only the non-redundant probe locations. Later, when the application is executed, the probes are used at a block 208 to collect information that includes, but is not necessarily limited to, information regarding execution of the probes' associated functions. In particular, this information may include, for example, the address of the function being entered or exited, the address from which another function is being called, and/or the address of a function being called. The collected information may also include a stack pointer to facilitate tracing of previous system activity, as well as a time stamp indicating when a function was entered or exited, as discussed above. After this information is collected, it is analyzed at a block 210 to determine how time should be attributed among the various functions making up the application.

By eliminating those probes that would produce redundant information, the various implementations of the present invention reduce the amount of data collection activity performed while analyzing the performance of an application, thereby reducing the extent to which the analyzed application's operation is distorted. As a result, the accuracy of the profiling process is improved.

While the embodiments of the invention have been described with specific focus on their embodiment in a software implementation, the invention as described above is not limited to software embodiments. For example, the invention may be implemented in whole or in part in hardware, firmware, software, or any combination thereof. The software of the invention may be embodied in various forms such as a computer program encoded in a machine-readable medium, such as a CD-ROM, magnetic medium, ROM or RAM, or in an electronic signal. Further, as used in the claims herein, the term "module" shall mean any hardware or software component, or any combination thereof.

What is claimed is:

1. A computer-implemented method for collecting information relating to
5 execution of an application, the method comprising:
determining a set of probe locations in the application at which collecting data
relating to the execution of the application would produce non-redundant information;
and
inserting probes only at the determined probe locations in the application.
10
2. The method of claim 1, further comprising determining entry and exit
points of a plurality of functions constituting at least a portion of the application.
3. The method of claim 2, further comprising identifying the entry and exit
15 points as probe locations at which probes are to be inserted.
4. The method of claim 1, further comprising:
identifying a first location within the application at which a function call directs
execution of the application to a second location outside of a current module; and
20 inserting a first probe before the identified first location and a second probe after
the identified first location.
5. The method of claim 4, wherein the first probe is configured to collect an
address of a first function in which the identified first location is located, an address of a
25 second function in which the second location is located, a first stack pointer, and a first
time indicator, and the second probe is configured to collect the address of the second
function, a second stack pointer, and a second time indicator.
6. The method of claim 1, further comprising:

identifying a first location within a calling function at which execution of the application is directed to a called function having an exit point at which execution of the application is directed to a second location outside of the calling function; and

5 inserting a first probe before the first location and a second probe after the second location.

7. The method of claim 6, wherein the first probe is configured to collect an address of the calling function, an address of the called function, a first stack pointer, and a first time indicator, and the second probe is configured to collect the address of the
10 called function, a second stack pointer, and a second time indicator.

8. The method of claim 1, further comprising:
identifying a block of code to which execution of the application is directed upon occurrence of an error; and
15 inserting a first probe at a beginning of the identified block of code and a second probe at an end of the identified block of code.

9. The method of claim 8, wherein the first probe is configured to collect an address of the block of code, a first stack pointer, and a first time indicator, and the
20 second probe is configured to collect the address of the block of code, a second stack pointer, and a second time indicator.

10. The method of claim 1, further comprising using the inserted probes to collect the information relating to the execution of the application.
25

11. The method of claim 10, further comprising analyzing the collected information.

12. A computer-implemented method for collecting information relating to
30 execution of an application, the method comprising:

determining entry and exit points of a plurality of functions constituting at least a portion of the application;

determining a set of probe locations in the application at which collecting data relating to the execution of the application would produce non-redundant information, the

5 set of probe locations including at least the entry and exit points of the functions;

inserting probes only at the determined probe locations in the application;

using the inserted probes to collect the information relating to the execution of the application; and

analyzing the collected information.

10

13. The method of claim 12, further comprising:

identifying a first location within the application at which a function call directs execution of the application to a second location outside of a current module; and

15 inserting a first probe before the identified first location and a second probe after the identified first location,

wherein the first probe is configured to collect an address of a first function in which the identified first location is located, an address of a second function in which the second location is located, a first stack pointer, and a first time indicator, and the second probe is configured to collect the address of the second function, a second stack pointer,

20

and a second time indicator.

14. The method of claim 12, further comprising:

identifying a first location within a calling function at which execution of the application is directed to a called function having an exit point at which execution of the application is directed to a second location outside of the calling function; and

25

inserting a first probe before the first location and a second probe after the second location,

wherein the first probe is configured to collect an address of the calling function, an address of the called function, a first stack pointer, and a first time indicator, and the second probe is configured to collect the address of the called function, a second stack pointer, and a second time indicator.

30

15. The method of claim 12, further comprising:
identifying a block of code to which execution of the application is directed upon occurrence of an error; and

5 inserting a first probe at a beginning of the identified block of code and a second probe at an end of the identified block of code,

wherein the first probe is configured to collect an address of the block of code, a first stack pointer, and a first time indicator, and the second probe is configured to collect the address of the block of code, a second stack pointer, and a second time indicator.

10 ~~16.~~ A computer-readable medium having computer-executable instructions for:

determining a set of probe locations in the application at which collecting data relating to the execution of the application would produce non-redundant information;

15 and

inserting probes only at the determined probe locations in the application.

17. The computer-readable medium of claim 16, having further computer-executable instructions for determining entry and exit points of a plurality of functions constituting at least a portion of the application.

18. The computer-readable medium of claim 17, having further computer-executable instructions for identifying the entry and exit points as probe locations at which probes are to be inserted.

19. The computer-readable medium of claim 16, having further computer-executable instructions for:

identifying a first location within the application at which a function call directs execution of the application to a second location outside of a current module; and

30 inserting a first probe before the identified first location and a second probe after the identified first location.

20. The computer-readable medium of claim 19, wherein the first probe is configured to collect an address of a first function in which the identified first location is located, an address of a second function in which the second location is located, a first stack pointer, and a first time indicator, and the second probe is configured to collect the address of the second function, a second stack pointer, and a second time indicator.

21. The computer-readable medium of claim 16, having further computer-executable instructions for:

identifying a first location within a calling function at which execution of the application is directed to a called function having an exit point at which execution of the application is directed to a second location outside of the calling function; and inserting a first probe before the first location and a second probe after the second location.

22. The computer-readable medium of claim 21, wherein the first probe is configured to collect an address of the calling function, an address of the called function, a first stack pointer, and a first time indicator, and the second probe is configured to collect the address of the called function, a second stack pointer, and a second time indicator.

23. The computer-readable medium of claim 16, having further computer-executable instructions for:

identifying a block of code to which execution of the application is directed upon occurrence of an error; and inserting a first probe at a beginning of the identified block of code and a second probe at an end of the identified block of code.

24. The computer-readable medium of claim 23, wherein the first probe is configured to collect an address of the block of code, a first stack pointer, and a first time

indicator, and the second probe is configured to collect the address of the block of code, a second stack pointer, and a second time indicator.

25. The computer-readable medium of claim 16, having further computer-executable instructions for using the inserted probes to collect the information relating to the execution of the application.

26. The computer-readable medium of claim 25, having further computer-executable instructions for analyzing the collected information.

27. A computer-readable medium having computer-executable instructions for:

- determining entry and exit points of a plurality of functions constituting at least a portion of the application;
- determining a set of probe locations in the application at which collecting data relating to the execution of the application would produce non-redundant information, the set of probe locations including at least the entry and exit points of the functions;
- inserting probes only at the determined probe locations in the application;
- using the inserted probes to collect the information relating to the execution of the application; and
- analyzing the collected information.

28. The computer-readable medium of claim 27, having further computer-executable instructions for:

- identifying a first location within the application at which a function call directs execution of the application to a second location outside of a current module; and

- inserting a first probe before the identified first location and a second probe after the identified first location,

- wherein the first probe is configured to collect an address of a first function in which the identified first location is located, an address of a second function in which the second location is located, a first stack pointer, and a first time indicator, and the second

probe is configured to collect the address of the second function, a second stack pointer, and a second time indicator.

29. The computer-readable medium of claim 27, having further computer-executable instructions for:

identifying a first location within a calling function at which execution of the application is directed to a called function having an exit point at which execution of the application is directed to a second location outside of the calling function; and

inserting a first probe before the first location and a second probe after the second location,

wherein the first probe is configured to collect an address of the calling function, an address of the called function, a first stack pointer, and a first time indicator, and the second probe is configured to collect the address of the called function, a second stack pointer, and a second time indicator.

30. The computer-readable medium of claim 27, having further computer-executable instructions for:

identifying a block of code to which execution of the application is directed upon occurrence of an error; and

inserting a first probe at a beginning of the identified block of code and a second probe at an end of the identified block of code,

wherein the first probe is configured to collect an address of the block of code, a first stack pointer, and a first time indicator, and the second probe is configured to collect the address of the block of code, a second stack pointer, and a second time indicator.

31. A computer arrangement configured to execute computer-executable instructions for:

determining a set of probe locations in the application at which collecting data relating to the execution of the application would produce non-redundant information;

and

inserting probes only at the determined probe locations in the application.

32. The computer arrangement of claim 31, further configured to execute computer-executable instructions for determining entry and exit points of a plurality of functions constituting at least a portion of the application.

5

33. The computer arrangement of claim 32, further configured to execute computer-executable instructions for identifying the entry and exit points as probe locations at which probes are to be inserted.

10

34. The computer arrangement of claim 31, further configured to execute computer-executable instructions for:

identifying a first location within the application at which a function call directs execution of the application to a second location outside of a current module; and

15

inserting a first probe before the identified first location and a second probe after the identified first location.

20

35. The computer arrangement of claim 34, wherein the first probe is configured to collect an address of a first function in which the identified first location is located, an address of a second function in which the second location is located, a first stack pointer, and a first time indicator, and the second probe is configured to collect the address of the second function, a second stack pointer, and a second time indicator.

25

36. The computer arrangement of claim 31, further configured to execute computer-executable instructions for:

identifying a first location within a calling function at which execution of the application is directed to a called function having an exit point at which execution of the application is directed to a second location outside of the calling function; and

inserting a first probe before the first location and a second probe after the second location.

30

37. The computer arrangement of claim 36, wherein the first probe is configured to collect an address of the calling function, an address of the called function, a first stack pointer, and a first time indicator, and the second probe is configured to collect the address of the called function, a second stack pointer, and a second time indicator.

38. The computer arrangement of claim 31, further configured to execute computer-executable instructions for:

identifying a block of code to which execution of the application is directed upon occurrence of an error; and

inserting a first probe at a beginning of the identified block of code and a second probe at an end of the identified block of code.

39. The computer arrangement of claim 38, wherein the first probe is configured to collect an address of the block of code, a first stack pointer, and a first time indicator, and the second probe is configured to collect the address of the block of code, a second stack pointer, and a second time indicator.

40. The computer arrangement of claim 31, further configured to execute computer-executable instructions for using the inserted probes to collect the information relating to the execution of the application.

41. The computer arrangement of claim 40, further configured to execute computer-executable instructions for analyzing the collected information.

42. A computer arrangement configured to execute computer-executable instructions for:

determining entry and exit points of a plurality of functions constituting at least a portion of the application;

determining a set of probe locations in the application at which collecting data relating to the execution of the application would produce non-redundant information, the set of probe locations including at least the entry and exit points of the functions;

inserting probes only at the determined probe locations in the application;

5 using the inserted probes to collect the information relating to the execution of the application; and

analyzing the collected information.

43. The computer arrangement of claim 42, further configured to execute
10 computer-executable instructions for:

identifying a first location within the application at which a function call directs execution of the application to a second location outside of a current module; and

inserting a first probe before the identified first location and a second probe after the identified first location,

15 wherein the first probe is configured to collect an address of a first function in which the identified first location is located, an address of a second function in which the second location is located, a first stack pointer, and a first time indicator, and the second probe is configured to collect the address of the second function, a second stack pointer, and a second time indicator.

20 44. The computer arrangement of claim 42, further configured to execute computer-executable instructions for:

identifying a first location within a calling function at which execution of the application is directed to a called function having an exit point at which execution of the
25 application is directed to a second location outside of the calling function; and

inserting a first probe before the first location and a second probe after the second location,

30 wherein the first probe is configured to collect an address of the calling function, an address of the called function, a first stack pointer, and a first time indicator, and the second probe is configured to collect the address of the called function, a second stack pointer, and a second time indicator.

ABSTRACT OF THE DISCLOSURE

A method, apparatus, and software are disclosed for collecting information relating to the execution of an application for use in analyzing the performance of the application. Probes are inserted at only those points in the application that would yield
5 non-redundant information. As a result, the impact of the information collection process on the application is reduced, improving the usefulness of the collected information.

"Express Mail" mailing label number: EL510128190US

Date of Deposit: Apr 1 26, 2000

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231

Printed Name SHAWN L. CHISE

Signature [Signature]

009240" 6920560




FIG. 1

Table 1. Demographic characteristics of the study population	
Age (years)	50.0 ± 10.0
Gender	
Male	50.0%
Female	50.0%
Education (years)	12.0 ± 2.0
Marital status	
Married	80.0%
Single	20.0%
Occupation	
Professional	30.0%
Managerial	20.0%
Technical	10.0%
Service	20.0%
Unemployed	20.0%
Income (USD/month)	1,000.0 ± 500.0
Health status	
Good	70.0%
Fair	20.0%
Poor	10.0%

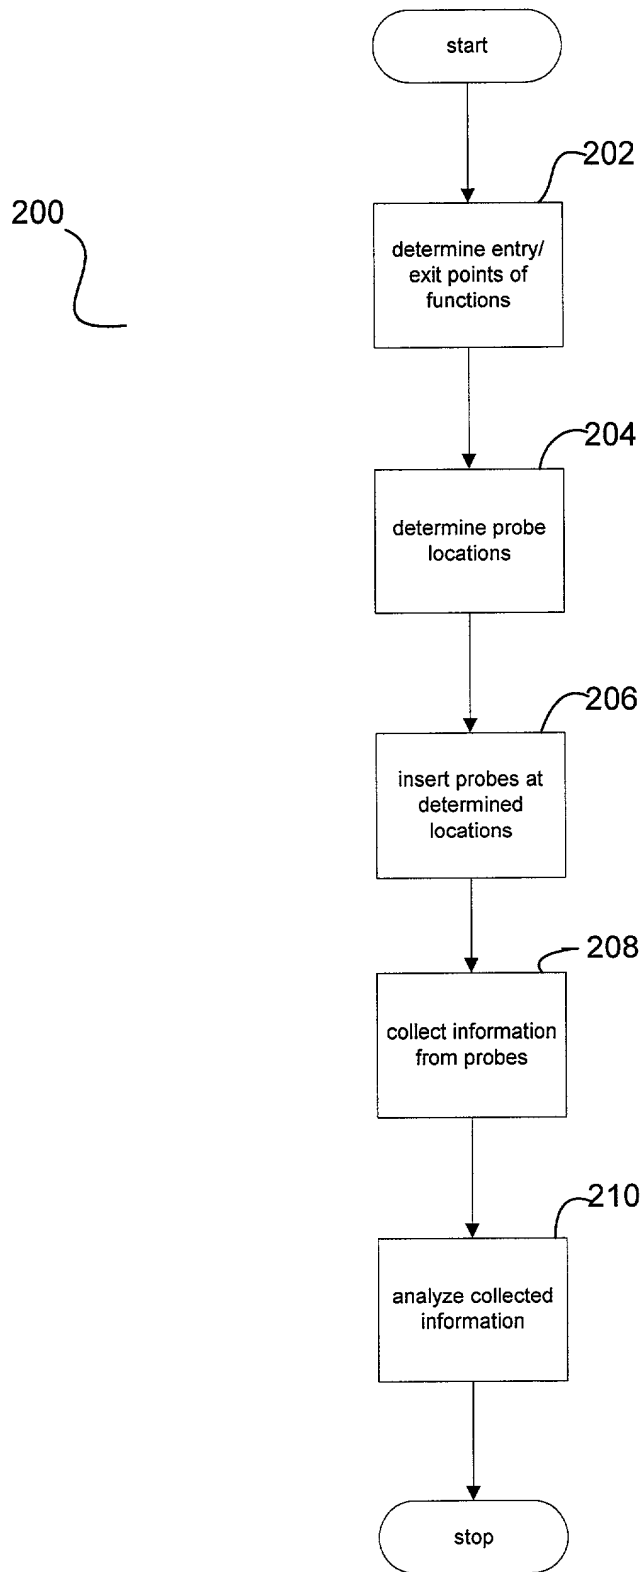


FIG. 2

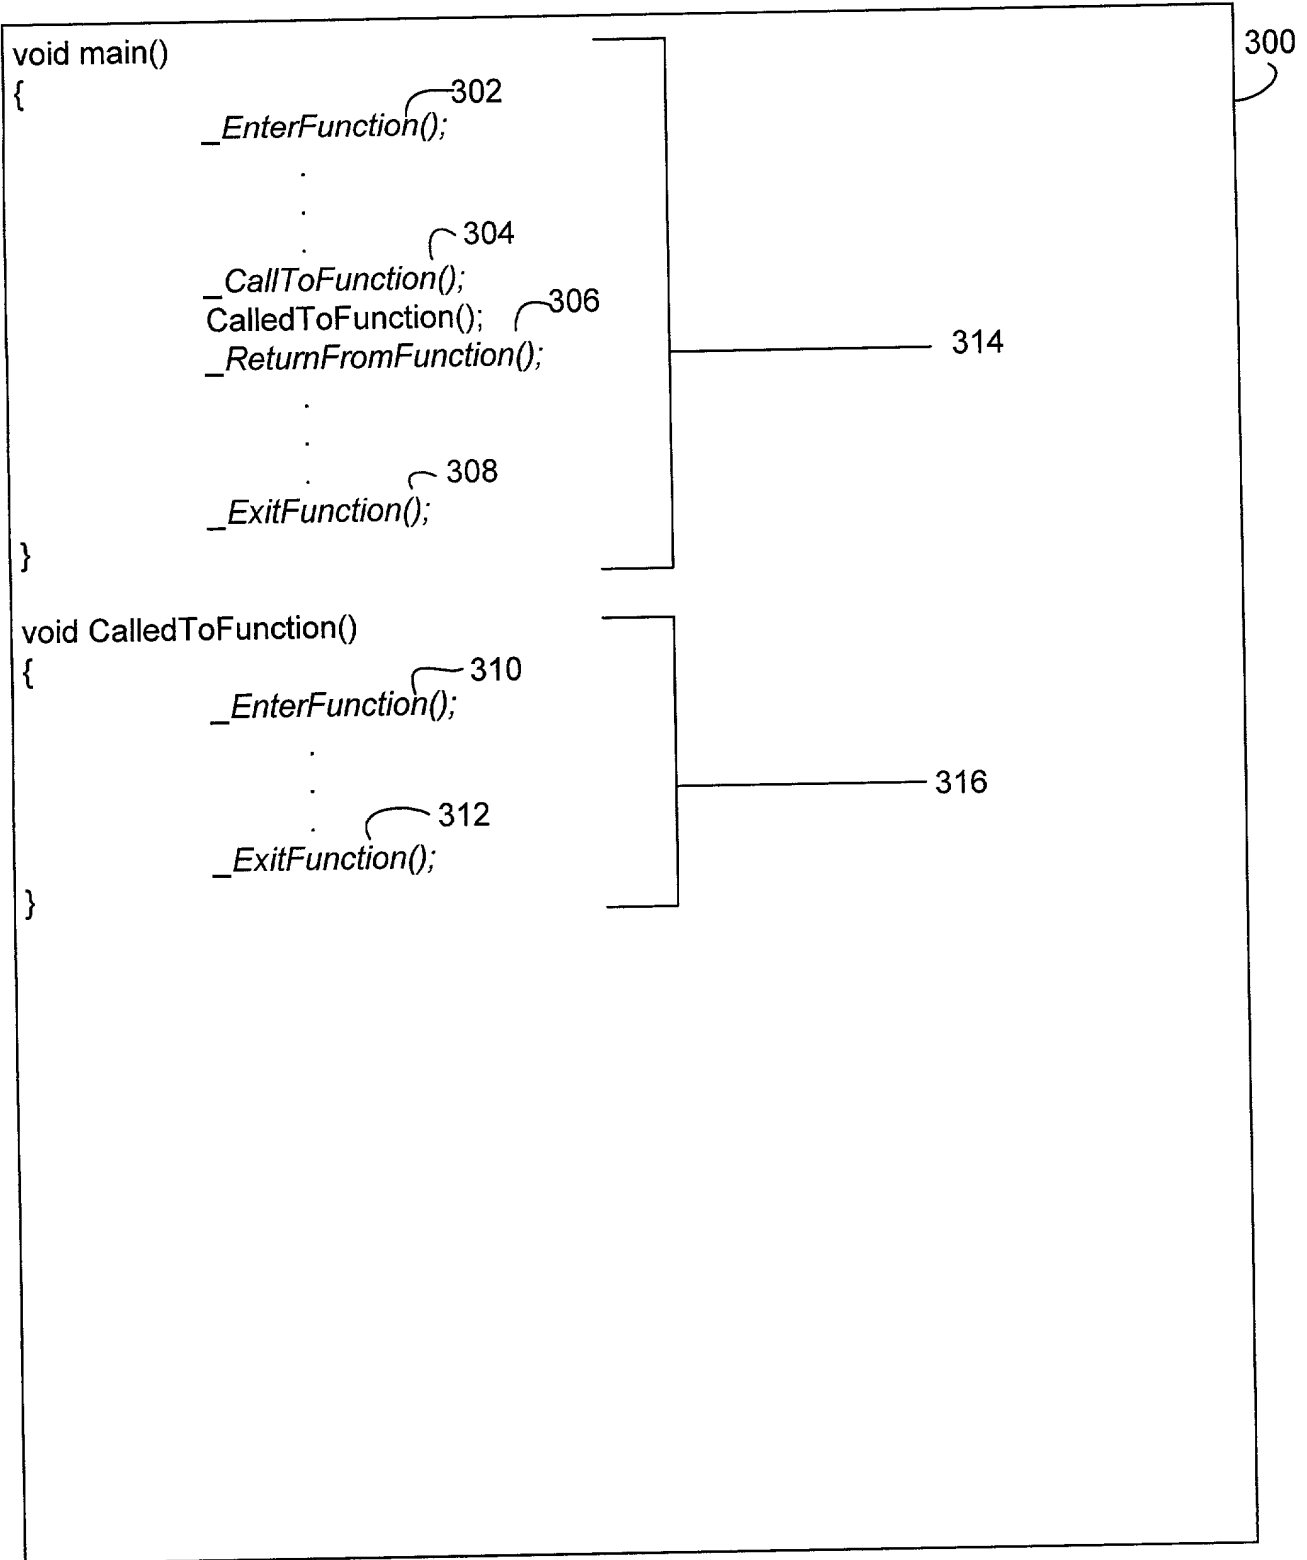


FIG. 3

```

void main()
{
    .
    .
    .
    call Function1();      416
    _ReturnFromFunction();
    .
    .
}

void Function1()
{
    .
    .
    .
    _CallToFunction();     412
    jump to Function2();   406
}

void Function2()
{
    .
    .
    .
    return;               410
    // this returns to main(), not to Function1()
}

```

Diagram annotations:

- 402: Bracket indicating the scope of the `main()` function.
- 414: Bracket indicating the scope of the call to `Function1()` and the subsequent `_ReturnFromFunction()` call.
- 404: Bracket indicating the scope of the `Function1()` function.
- 408: Bracket indicating the scope of the `Function2()` function.

FIG. 4

SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.

United States Patent Application

COMBINED DECLARATION AND POWER OF ATTORNEY

As a below named inventor I hereby declare that: my residence, post office address and citizenship are as stated below next to my name; that

I verily believe I am the original, first and sole inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled: **MINIMAL INSTRUMENTATION FOR LOSSLESS CALL PROFILING**.

The specification of which is attached hereto.

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with 37 C.F.R. § 1.56 (attached hereto). I also acknowledge my duty to disclose all information known to be material to patentability which became available between a filing date of a prior application and the national or PCT international filing date in the event this is a Continuation-In-Part application in accordance with 37 C.F.R. § 1.63(e).

I hereby claim foreign priority benefits under 35 U.S.C. § 119(a)-(d) or 365(b) of any foreign application(s) for patent or inventor's certificate, or 365(a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on the basis of which priority is claimed:

No such claim for priority is being made at this time.

I hereby claim the benefit under 35 U.S.C. § 119(e) of any United States provisional application(s) listed below:

No such claim for priority is being made at this time.

I hereby claim the benefit under 35 U.S.C. § 120 or 365(c) of any United States and PCT international application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT international application in the manner provided by the first paragraph of 35 U.S.C. § 112, I acknowledge the duty to disclose material information as defined in 37 C.F.R. § 1.56(a) which became available between the filing date of the prior application and the national or PCT international filing date of this application:

No such claim for priority is being made at this time.

I hereby appoint the following attorney(s) and/or patent agent(s) to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith:

Anglin, J. Michael	Reg. No. 24,916	Jurkovich, Patti J.	Reg. No. 44,813	Nelson, Albin J.	Reg. No. 28,650
Bianchi, Timothy E.	Reg. No. 39,610	Kalis, Janal M.	Reg. No. 37,650	Nielsen, Walter W.	Reg. No. 25,539
Billion, Richard E.	Reg. No. 32,836	Kaufmann, John D.	Reg. No. 24,017	Oh, Allen J.	Reg. No. 42,047
Black, David W.	Reg. No. 42,331	Klima-Silberg, Catherine I.	Reg. No. 40,052	Padys, Danny J.	Reg. No. 35,635
Brennan, Leoniede M.	Reg. No. 35,832	Kluth, Daniel J.	Reg. No. 32,146	Parker, J. Kevin	Reg. No. 33,024
Brennan, Thomas F.	Reg. No. 35,075	Lacy, Rodney L.	Reg. No. 41,136	Perdok, Monique M.	Reg. No. 42,989
Brooks, Edward J., III	Reg. No. 40,925	Leffert, Thomas W.	Reg. No. 40,697	Prout, William F.	Reg. No. 33,995
Chu, Dinh C.P.	Reg. No. 41,676	Lemaire, Charles A.	Reg. No. 36,198	Sako, Katie E.	Reg. No. 32,628
Clark, Barbara J.	Reg. No. 38,107	Litman, Mark A.	Reg. No. 26,390	Schumm, Sherry W.	Reg. No. 39,422
Crouse, Daniel D.	Reg. No. 32,022	Lundberg, Steven W.	Reg. No. 30,568	Schwegman, Micheal L.	Reg. No. 25,816
Dahl, John M.	Reg. No. 44,639	Mack, Lisa K.	Reg. No. 42,825	Smith, Michael G.	Reg. No. 45,368
Drake, Eduardo E.	Reg. No. 40,594	Maeyaert, Paul L.	Reg. No. 40,076	Speier, Gary J.	Reg. No. 45,458
Eliseeva, Maria M.	Reg. No. 43,328	Maki, Peter C.	Reg. No. 42,832	Steffey, Charles E.	Reg. No. 25,179
Embretson, Janet E.	Reg. No. 39,665	Malen, Peter L.	Reg. No. 44,894	Terry, Kathleen R.	Reg. No. 31,884
Fordenbacher, Paul J.	Reg. No. 42,546	Mates, Robert E.	Reg. No. 35,271	Tong, Viet V.	Reg. No. 45,416
Forrest, Bradley A.	Reg. No. 30,837	McCrackin, Ann M.	Reg. No. 42,858	Viksnins, Ann S.	Reg. No. 37,748
Harris, Robert J.	Reg. No. 37,346	Nama, Kash	Reg. No. 44,255	Woessner, Warren D.	Reg. No. 30,440
Huebsch, Joseph C.	Reg. No. 42,673				

I hereby authorize them to act and rely on instructions from and communicate directly with the person/assignee/attorney/firm/organization/who/which first sends/sent this case to them and by whom/which I hereby declare that I have consented after full disclosure to be represented unless/until I instruct Schwegman, Lundberg, Woessner & Kluth, P.A. to the contrary.

Please direct all correspondence in this case to **Schwegman, Lundberg, Woessner & Kluth, P.A.** at the address indicated below:
P.O. Box 2938, Minneapolis, MN 55402
Telephone No. (612)373-6900

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of sole inventor : **Barry M. Nolte**
Citizenship: **United States of America** Residence: **Kirkland, WA**
Post Office Address: **13123 NE 111th Place**
Kirkland, WA 98033

Signature: _____ Date: _____
Barry M. Nolte

Full Name of inventor:
Citizenship: Residence:
Post Office Address:

Signature: _____ Date: _____

§ 1.56 Duty to disclose information material to patentability.

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclose information exists with respect to each pending claim until the claim is canceled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is canceled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclose all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§ 1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:

- (1) prior art cited in search reports of a foreign patent office in a counterpart application, and
- (2) the closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.

(b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made of record in the application, and

- (1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or
- (2) It refutes, or is inconsistent with, a position the applicant takes in:
 - (i) Opposing an argument of unpatentability relied on by the Office, or
 - (ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

- (1) Each inventor named in the application;
- (2) Each attorney or agent who prepares or prosecutes the application; and
- (3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.

(d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent, or inventor.